Analysis of the IJCNN 2011 UTL challenge

Isabelle Guyon[a], Gideon Dror[b], Vincent Lemaire[c], Daniel L. Silver[d],
Graham Taylor[e], and David W. Aha[f]

*a. Clopinet, California, USA. b. Yahoo!, Haifa, Israel. c. Orange Labs, France. d. Acadia
University, Canada. e. New York University, USA. f. Naval Research Laboratory*

Abstract

We organized a challenge in "Unsupervised and Transfer Learning": the UTL challenge (`http://clopinet.com/ul`). We made available large datasets from various application domains: handwriting recognition, image recognition, video processing, text processing, and ecology. The goal was to learn data representations that capture regularities of an input space for re-use across tasks. The representations were evaluated on supervised learning "target tasks" unknown to the participants. The first phase of the challenge was dedicated to "unsupervised transfer learning" (the competitors were given only unlabeled data). The second phase was dedicated to "cross-task transfer learning" (the competitors were provided with a limited amount of labeled data from "source tasks", distinct from the "target tasks"). The analysis indicates that learned data representations yield significantly better results than those obtained with original data or data preprocessed with standard normalizations and functional transforms.

Keywords: machine learning, transfer learning, unsupervised learning, metric learning, kernel learning, unlabeled data, challenge, competition

1. Introduction

Classical machine learning techniques, including artificial neural networks, work well if data are independently and identically distributed and if there there are sufficient numbers of labeled training data. Unfortunately, in many real world applications, these assumptions are violated: (1) Data available for training are not always similar to data the system will be exposed to when it is deployed. (2) Few labeled training data may be available due to the cost or burden of manually annotating data. Transfer learning addresses both of these shortcomings by providing tools to learn a new task for which labeled data are scarce from a related task for which data are abundant, labeled or not (Figure 1).

Report Documentation Page

1. REPORT DATE **13 JAN 2012**	2. REPORT TYPE	3. DATES COVERED **00-00-2012 to 00-00-2012**

4. TITLE AND SUBTITLE **Analysis of the IJCNN 2011 UTL challenge**	5a. CONTRACT NUMBER
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Clopinet,955 Creston road ,Berkeley,CA,94708**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
Presented at the IEEE International Joint Conference on Neural Networks, San Jose, California, July 31 - August 5, 2011

14. ABSTRACT

We organized a challenge in ?Unsupervised and Transfer Learning?: the UTL challenge (http //clopinet.com/ul). We made available large datasets from various application domains handwriting recognition, image recognition, video processing, text processing, and ecology. The goal was to learn data representations that capture regularities of an input space for re-use across tasks. The representations were evaluated on supervised learning ?target tasks? unknown to the participants. The first phase of the challenge was dedicated to ?unsupervised transfer learning? (the competitors were given only unlabeled data). The second phase was dedicated to ?crosstask transfer learning? (the competitors were provided with a limited amount of labeled data from ?source tasks?, distinct from the ?target tasks?). The analysis indicates that learned data representations yield significantly better results than those obtained with original data or data preprocessed with standard normalizations and functional transforms.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**12**	

Figure 1: The Transfer Learning question: Can learning about house cats help us learn about tigers?

2. Definition of Unsupervised and Transfer Learning

In this paper, we call "domain" the input space (*e.g.,* a feature vector space) and we call "task" the output space (represented by labels for classification problems). We use the adjective "source" for an auxiliary problem, for which we have an abundance of data (*e.g.,* classifying house cats), and "target" for the problem of interest (*e.g.,* classifying tigers). A transfer learning system recognizes and applies knowledge and skills learned on a "source" problem to a novel "target" problem. Within this framework, there several settings, depending on whether **labels are available for the source task and/or the target task**, and **whether domains and tasks are the same or different for the source and the target problem** [1]. Some familiar settings include:

Inductive Transfer Learning: Labeled data are available for the target tasks to perform supervised learning, but not in abundance. A typical setting is that of a large multi-class problem, in which some classes are a lot more depleted than others (*e.g.,* in image classification or text classification). The most abundant classes serve as "source tasks" for the least abundant ones, representing the "target tasks". If all source data are labeled, one talks of **multi-task learning**. When unlabeled source data are available, the term **self-taught learning** [2] is used.

Domain Adaptation: Labeled data are available for both the source and the target tasks, but the input distribution (domain) is changing. For instance a speech recognition system would be trained with native speakers and be used by a non-native speaker willing to partially re-train it.

Semi-supervised and Transductive Learning: There is no change in either domain or task and labeled training data are available. In addition, either unlabeled training data are available (**semi-supervised learning**) or the unlabeled test data may be used for learning (**transductive learning**). An example of semi-supervised learning would be classifying images of skin lesions into cancer and non-cancer from a database of example including few confirmed cases.

2

Unsupervised Learning and Cross-task Transfer Learning: When no labeled data are available at all (in the source or target domain), **unsupervised learning** techniques (clustering, vector quantization, factor analysis, manifold learning, etc.) may be applied in an effort to represent data in a simpler and/or better way, either for visualization, or as data preprocessing, for future use in supervised learning tasks. When a limited amount of labeled data is available *only* in the source domain, we talk about **cross-task transfer learning**.

3. Setting and tasks of the challenge

The case examined in this challenge was that of Unsupervised and Cross-Task Transfer Learning. Labels were available for target tasks to the challenge organizers only, to evaluate data representations provided by the participants. This setting may seem artificial because there are ultimately labels available for some target tasks, so why not make same available to the participants? Our goal was to decouple problems and see how far one could get by working only on learning data representations. This exercise has turned out to be extremely fruitful.

The datasets of the challenge (Table 1) were split into a large development set[1], a validation set and a final evaluation set. **The goal of the challenge was to produce good data representations on the final evaluation set.** The validation set is similar to the final evaluation set; it was provided for practice. The assessment of the data representations was carried out on *target tasks* (that are supervised learning classification tasks), using labels known only to the competition organizers. The target tasks for the validation set and the final evaluation set are different but related. The intention is to determine the extent to which the abstract features are useful for classifying a family of related tasks. During the development period, online feedback was provided only on the validation set. The results on the final evaluation set were revealed only at the end of the challenge.

The challenge proceeded in two phases. The first phase focused on **unsupervised learning**.[2] During that phase, no labels were provided to the participants in either the source or the target domain. It was then followed by a second phase

---

[1]We use the nomenclature "development set" rather than "training set" to stress that the actual (supervised) training is performed with labeled target domain data by the organizers only. The development set is not the training set for the target tasks.

[2]From the point of view of the overall problem, including the evaluation on supervised target tasks performed by the organizers, this setting is similar to self-taught learning [2] because of the availability of unlabeled data in a source domain distinct from the target domain.

on **cross-task transfer learning** for which some labels for *source tasks, distinct from the target tasks*, were provided for a subset of the development data.

Table 1: UTL challenge datasets

*D*evel = number of examples in development data. *T*ransf = number of source task labels released in the second phase. The validation and final evaluation sets consist of 4096 examples each.

Dataset	Domain	Features	Sparsity	Devel.	Transf.
AVICENNA	Handwriting	120	0%	150205	50000
HARRY	Video	5000	98.1%	69652	20000
RITA	Images	7200	1.1%	111808	24000
SYLVESTER	Ecology	100	0%	572820	100000
TERRY	Text	47236	99.8%	217034	40000

We selected five different application domains that are illustrative of fields in which transfer learning is applicable:[3]

AVICENNA: Professor Mohamed Chériet, École de Technologie Supérieure, University of Quebec, Montréal, Canada, and his students prepared a large corpus of historical Arabic documents [3]. Transfer learning methods could accelerate the application of handwriting recognizers to historical manuscript by reducing the need for using human experts to labels and index them.

HARRY: The identification and recognition of gestures, postures and human behaviors has gained importance in applications such as video surveillance, gaming, marketing, computer interfaces and interpretation of sign languages for the deaf. The HARRY dataset was constructed from the KTH human action recognition dataset of Ivan Laptev and Barbara Caputo[4] and the Hollywood 2 dataset of human actions and scenes of Marcin Marszalek, Ivan Laptev, and Cordelia Schmid[5].

RITA: Object recognition in images is a classical pattern recognition task gaining importance for retrieval applications, including for Internet search. The RITA dataset was constructed from the CIFAR dataset of Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton[6], a subset of the 80 million Tiny images dataset of

---

[3]A detailed technical report on the datasets was made available after the challenge ended: `http://www.causality.inf.ethz.ch/ul_data/DatasetsUTLChallenge.pdf`

[4]`http://www.nada.kth.se/cvap/actions/`

[5]`http://www.irisa.fr/vista/Equipe/People/Laptev/download.html`

[6]`http://www.cs.toronto.edu/~kriz/cifar.html`

Antonio Torralba, Rob Fergus, and William T. Freeman[7].

SYLVESTER: Massive datasets need to be processed automatically to assist experts in ecology, geography, geology, climatology, archaeology, and seismology. We used data from the US Forest service to illustrate such tasks.[8]

TERRY: Internet search engines process billions of queries daily to rank web pages. Very few labeled data are available, but millions of documents must be indexed. We used a subset of the RCV1-v2 Text Categorization Test Collection derived from Reuter's news articles formatted and made publicly available by David Lewis [4].

These are all multi-class problems. We selected a subset of the classes as target tasks and used the remainder as source tasks. The tasks in the source and target domains are very distinct, for instance, the source domain may include pictures of cars, houses, trees, and the target domain would include horses, chairs, and computers. The datasets varied in difficulty, as illustrated by the histograms of performances of the participants (Figure 2). AVICENNA is very hard, HARRY and TERRY are the easiest tasks (but there is a very wide range of results) and RITA and SYLVESTER are of medium difficulty.

4. Protocol and evaluation

The challenge protocol was inspired by previous competitions we organized [5] and was designed to ensure fairness of the evaluation and stimulate participation. We provided guidance to the participants with detailed answers to Frequently Asked Questions (FAQ)[9] and we posted a short video tutorial on transfer learning.[10] The rules can be found on the website of the challenge.[11]

The data representations were assessed automatically on the website of the challenge. The evaluation software and sample code were provided to the participants. The organizers defined several binary classification tasks unknown to the participants and placed them in either the validation set or the final evaluation set. The platform used the data representations provided by the participants to train a linear classifier to solve these tasks. So in all cases, the data representa-

---

[7]http://groups.csail.mit.edu/vision/TinyImages/

[8]http://archive.ics.uci.edu/ml/datasets/Covertype

[9]http://www.causality.inf.ethz.ch/unsupervised-learning.php?page=FAQ

[10]http://www.youtube.com/watch?v=9ChVn3xVNDI
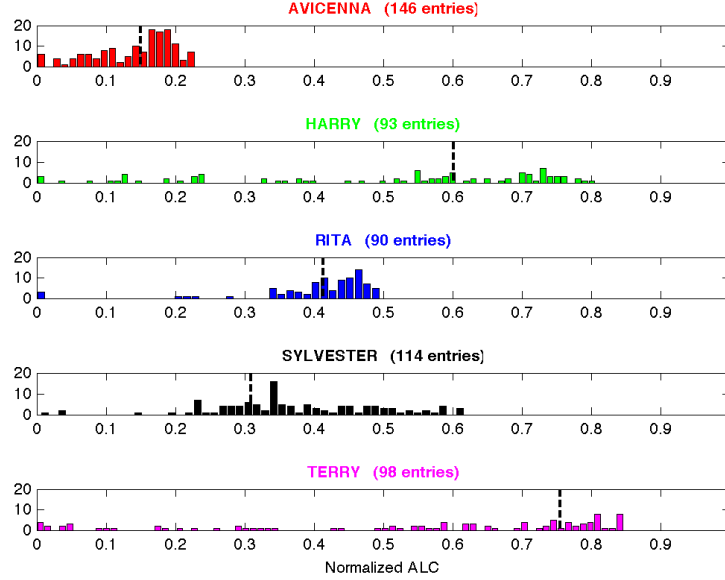
[11]http://clopinet.com/ul

Figure 2: Histogram of participant results for the first phase of the challenge. The y axis is the number of entries in each bin. Similar histograms were obtained in the second phase (not shown)

tions were assessed with supervised learning tasks. We chose a simple Hebbian learning rule.[12] The Hebbian algorithm has a kernelized version. Thus, we let the participants submit either a data representation or a semi-definite positive matrix (interpreted as a matrix of similarity between pairs of examples), which was used as a kernel. See [6] for details.

To compute the ranking score, a form of cross-validation was performed by partitioning randomly the evaluation data (validation set or final evaluation set) multiple times into a training and a test set, and averaging performances. The number of training examples m was varied from 1 to 2^6 and the area under the ROC curve[13] (AUC) was plotted against m in a log2 scale to emphasize the results on small m. The area under the learning curve (ALC) was used as a scoring metric to assess the results. The ALC criterion is a good way to aggregate the values of

---

[12]Experiments conducted after the challenge using the submissions of the top ranking participants revealed that other linear classifiers, not making an independence assumption between examples, do not yield performance improvements. This can be explained by the fact that it is possible to orthogonalize the examples with *e.g.,* PCA, making Hebbian learning quasi-optimal. Additionally, for small training sets, Hebbian learning is robust against overfitting.

[13]This is the area under the curve plotting hit rate *vs.* false alarm rate.

6

the AUC over all considered number of training examples. The participants were ranked by ALC for each individual dataset. The participants having submitted a complete experiment (*i.e.,* reporting results on all 5 datasets of the challenge) could enter the final ranking. The winner was determined by the best average rank over all datasets for the results on a complete experiment or their choice. See [6] for details.

5. Results

The challenge attracted 76 participants. There was more participation in the first phase than in the second phase: In the first phase, 6933 jobs were submitted, including 41 complete final entries, while, in the second phase, 1141 jobs were submitted including 14 complete final entries. There were in the end 16 ranked teams in the first phase and 8 ranked teams in the second phase. Not all teams decided to enter the final ranking, despite the option to preserve their anonymity.

The results of the top ranking teams are shown in Table 2. The "normalized" ALC refers to $(ALC - Arand)/(Amax - Arand)$, where Amax is the ALC obtained when perfect predictions are made and Arand is (in expectation) the ALC obtained for random predictions. We show in boldface the best result for both phases. The numbers in parentheses are the ranks for the individual datasets in each phase. The complete result tables are available online[14] and details are provided in [6]. In support of the significance of the results of the challenge, the top ranking teams in both phases used consistently the same principled methods on all datasets, and performed well on all of them.

It was important for us to assess whether unsupervised learning helps compared to classical normalizations or no preprocessing at all. We ran a couple of baseline algorithms for comparison. Using unsupervised learning, the participants outperformed the organizers on 4/5 datasets (for HARRY, the normalized representation achieved the best results in phase 1).

Finally, we examined the correlation between validation set and final evaluation set performances (Figure 3). The graph reveals that on several datasets the test set was easier that the validation set. We did this on purpose so the participants would not be frustrated. After removing a few outliers (probably due to submission errors) we obtained a correlation coefficient of 0.88 in the first phase and 0.89 in the second phase. Most participants simply used the validation set performance

---

[14]Result tables: `http://www.causality.inf.ethz.ch/ unsupervised-learning.php?page=results`

Table 2: Normalized ALC values of the top ranking participants.

Phase 1 - Unsupervised Learning							
Rank	Team	Experiment	Avicenna	Harry	Rita	Sylvester	Terry
1	AIO	AIO	0.2183 (1)	0.7043 (6)	0.4951 (1)	0.4569 (6)	**0.8465** (1)
2	1055A	exp1	0.1906 (6)	0.7357 (3)	0.4782 (5)	0.5828 (1)	0.8437 (2)
3	Airbus	A3XX	0.2174 (2)	0.7545 (2)	0.4724 (7)	0.4949 (4)	0.8390 (3)
4	LISA	LISA	0.1960 (5)	0.8062 (1)	0.4731 (6)	0.4763 (5)	0.7959 (6)
Phase 2 - Transfer Learning							
Rank	Team	Experiment	Avicenna	Harry	Rita	Sylvester	Terry
1	LISA	agartha	**0.2273** (1)	**0.8619** (1)	**0.5029** (1)	0.5650 (3)	0.8160 (2)
2	tkgw	crush	0.1973 (2)	0.7533 (2)	0.4095 (4)	**0.5933** (1)	0.8118 (3)
3	1055A	phase2exp1	0.1511 (4)	0.7381 (3)	0.4992 (2)	0.5873 (2)	0.8437 (1)
4	FAST	teaf	0.1909 (3)	0.3580 (4)	0.4275 (3)	0.3379 (5)	0.6485 (4)

as a model selection criterion. Because of the high correlation between validation performance and test performance, this turned out to be an effective strategy.

We surveyed the participants to determine what algorithms, software and hardware was used and the best entrants were asked to produce a full length paper, which will appear in JMLR W&CP [7]. We briefly summarize the methods and findings.

In the first phase, the winner (team name: AIO) used an algorithm to train kernels [8]. Using validation data, they incrementally improved their kernel, each time checking the performance on the leaderboard. They developed a systematic method of sequential kernel transformations and recorded which sequence ended up giving the best performance on validation data. Then, they applied the same sequence to the final evaluation data. The method can be interpreted as a greedy search for hyperparameters of a compound kernel. The AIO team did not use the development dataset for training at all: The cascade of steps selected with the validation set was then applied to the final evaluation data to produce a new kernel. Thus this method can be qualified of "transductive".

The team LISA, ranking first in the second phase and fourth in the first phase, based their solution on Deep Learning techniques, in particular for unsupervised learning of representations [9]. Their methods follow the techniques described in [10]. Those exploit as building blocks unsupervised learning of single-layer models, such as Restricted Boltzmann Machines, to construct deeper models such as Deep Belief Networks. Their last layer is a transductive PCA (PCA computed

Figure 3: Correlation between results on the validation and test data in phase 1. A similar plot was obtained with phase 2 data.

directly on the final evaluation data, for the feature representation produced by the preprocessor trained with the development and validation data).

The team 1055A, ranking second in the first phase and third in the second phase used classical unsupervised learning methods: Principal component analysis (PCA) and k-means clustering [11]. They first computed the principal components on validation dataset and used the on-line feedback to determine the first n principal components that gave the best global score. Clustering was then performed in the PCA representation and repeated 100 times with different class seeds. The number of clusters was optimized with the feedback from the validation dataset. They submitted data representations in a binary encoding of cluster membership.

The team tkgw, ranking second in the second phase, used a method called "Random Forest Proximity", which recursively searches for principal directions when going down a decision tree (halting at a depth of 12). Random Forests [12] are ensembles of decision trees built by resampling variables and training examples. The method allowed the authors to create a large number of features, from which they generated a similarity measure. The similarity measure was then turned into a semi-definite positive kernel matrix with a suitable normalization.

9

See the fact sheet for details.

The team Airbus, ranking third in the first phase, tried various preprocessing methods and selected the best one using the validation set. On AVICENNA they ended up using PCA with 90% of variance, and then a RBF kernel. For HARRY and TERRY, they rotated the representation then used a linear kernel. For RITA, they ran k-means to get some clusters then Maximum Variance Unfolding (MVU) on each of them to select features, and then a denoising algorithm. For SYLVESTER, they used whitening, taking 90% of the variance.

The setting of the challenge revealed the power of unsupervised learning as a preprocessor. For all the datasets, unsupervised learning produced results significantly better than the baseline methods (raw data or simple normalizations). The participants exploited effectively the feed-back received on the validation set to select the best data representations. The skepticism around the effectiveness of unsupervised learning is justified when no performance on a supervised task is available. However, unsupervised learning can be the object of model selection using a supervised task, similarly to preprocessing, feature selection, and hyper-parameter selection. An interesting new outcome of this challenge is that the supervised tasks used for model selection can be distinct from the tasks used for the final evaluation. So, even though the learning algorithms are unsupervised, transfer learning is happening at the model selection level.

In phase 1, there was a danger of overfitting by trying too many methods and relying too heavily on the performance on the validation set. One team for instance overfitted in phase 1, ranking 1st on the validation set, but only 4th on the final evaluation set. Possibly, criteria involving both the reconstruction error and the classification accuracy on the validation tasks may be more effective for model selection. This should be the subject of further research. In phase 2, the participants had available "transfer labels" for a subset of the development data (for classification tasks distinct from the classification tasks of the validation set and the final evaluation set). Therefore, they had the opportunity to use such labels to devise transfer learning strategies. The most effective strategy seems to have been to use the transfer labels for model selection again. None of the participants used those labels for learning.

6. Conclusions

The results of the challenge confirm results reported in the literature that unsupervised learning can be beneficial for preprocessing. The benefits of transfer learning in the "cross-task transfer learning" setting studied in this challenge are

mainly derived from an improvement in model selection of unsupervised learning preprocessing techniques. In particular, the challenge demonstrated that the validation data needs not to be drawn from the same distribution, it suffices that the source task used bears some resemblance to the target task. The importance of the degree of resemblance of the two tasks remains to be determined.

Overall, an array of algorithms were used, including classical linear methods like Principal Component Analysis (PCA), and non-linear methods like clustering (k-means and hierarchical clustering being the most popular), and Kernel-PCA (KPCA), as well as cutting edge methods including non-linear auto-encoders and restricted Bolzmann machines (RBMs) [9]. A general methodology seems to have emerged. Most top ranking participants used simple normalizations (like variable standardization and/or data sphering using PCA) as a first step, followed by one or several layers of non-linear processing (stacks of auto-encoders, RBMs, KPCA, and/or clustering). Finally, "transduction" played a key role in winning first place: either the whole preprocessing chain was applied directly to the final evaluation data (this is the strategy of Fabio Aiolli who won first place in phase 1 [8]); or alternatively, the final evaluation data, preprocessed with a preprocessor trained on development+validation data, was post-processed with PCA (so-called "transductive PCA" used by the LISA team, who won the second phase [9]).

Acknowledgments

References

[1] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Transactions on Knoweledge and Data Engineering*, vol. 22, no. 10, pp. 1345–1359, October 2010.

[2] R. Raina, A. Battle, H. Lee, B. Packer, and A. Y. Ng, "Self-taught learning: Transfer learning from unlabeled data," in *Proceedings of the Twenty-fourth International Conference on Machine Learning*, 2007.

[3] R. F. Moghaddam, M. Cheriet, M. M. Adankon, K. Filonenko, and R. Wisnovsky, "IBN SINA: a database for research on processing and understanding of arabic manuscripts images," in *Document Analysis Systems*, 2010, pp. 11–18.

[4] D. D. Lewis, Y. Yang, T. G. Rose, F. Li, G. Dietterich, and F. Li, "Rcv1: A new benchmark collection for text categorization research," *Journal of Machine Learning Research*, vol. 5, pp. 361–397, 2004.

[5] ChaLearn, "Challenges in machine learning." [Online]. Available: http://www.chalearn.org/challenges.html

[6] I. Guyon, G. Dror, V. Lemaire, G. Taylor, and A. David W, "Unsupervised and transfer learning challenge," in *The 2011 International Join Conference on Neural Networks*, July 2011, pp. 793–800.

[7] I. Guyon et al, Ed., *Unsupervised and Transfer Learning Workshop, in conjunction with ICML 2011*. JMLR W&CP, http://jmlr.csail.mit.edu/proceedings/papers, 2012, in press.

[8] F. Aiolli, "Transfer learning by kernel meta-learning," in *Unsupervised and Transfer Learning Workshop, in conjunction with ICML 2011*. JMLR W&CP, 2012, in press.

[9] G. Mesnil et al, "Unsupervised and transfer learning challenge: a deep learning approach," in *Unsupervised and Transfer Learning Workshop, in conjunction with ICML 2011*. JMLR W&CP, 2012, in press.

[10] Y. Bengio, "Learning deep architectures for AI," *Foundations and Trends in Machine Learning*, vol. 2, no. 1, pp. 1–127, 2009, also published as a book. Now Publishers, 2009.

[11] C. Liu, J. Xie, H. Xiong, and Y. Ge, "Stochastic unsupervised learning on unlabeled data," in *Unsupervised and Transfer Learning Workshop, in conjunction with ICML 2011*. JMLR W&CP, 2012, in press.

[12] L. Breiman, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.